

A Distance-Based Test of Association Between Paired Heterogeneous Genomic Data

Christopher Minas^{1,2}, Edward Curry³, and Giovanni Montana^{*1}

¹Statistics Section, Department of Mathematics, South Kensington Campus, Imperial College London, London, SW7 2AZ

²Institute of Clinical Sciences, Hammersmith Campus, Imperial College London, London, W12 0NN

³Ovarian Cancer Action Research Centre, Hammersmith Campus, Imperial College London, London, W12 0NN

March 29, 2013

Abstract

Due to rapid technological advances, a wide range of different measurements can be obtained from a given biological sample including single nucleotide polymorphisms, copy number variation, gene expression levels, DNA methylation and proteomic profiles. Each of these distinct measurements provides the means to characterize a certain aspect of biological diversity, and a fundamental problem of broad interest concerns the discovery of shared patterns of variation across different data types. Such data types are heterogeneous in the sense that they represent measurements taken at very different scales or described by very different data structures. We propose a distance-based statistical test, the generalized RV (GRV) test, to assess whether there is a common and non-random pattern of variability between paired biological measurements obtained from the same random sample. The measurements enter the test through distance measures which can be chosen to capture particular aspects of the data. An approximate null distribution is proposed to compute p-values in closed-form and without the need to perform costly Monte Carlo permutation procedures. Compared to the classical Mantel test for association between distance matrices, the GRV test has been found to be more powerful in a number of simulation settings. We also report on an application of the GRV test to detect biological pathways in which genetic variability is associated to variation in gene expression levels in ovarian cancer samples, and present results obtained from two independent cohorts.

1 Introduction

A proliferation in the development and application of high-throughput measurement platforms in biological research has resulted in the increasing availability of multiple levels of molecular data available for the same biological sampling units. For a range of human tumour types, for instance, the Cancer Genome Atlas (TCGA) consortium have made available

*Corresponding author: g.montana@imperial.ac.uk

genome-wide SNP genotypes, DNA copy number, genome-wide gene (and mRNA) expression, DNA methylation and proteomic profiles, as well as a range of mutation calls arising from deep sequencing. Such comprehensive molecular profiling of biological sampling unit provides opportunities for gaining insight into the mechanisms by which different molecular entities interact with one another to influence the overall state of the cell(s) in question. This is possible through comparing the different types of observations on each biological sampling unit, via some means of paired analysis of the individual datasets. A certain degree of complexity is inherent in this type of analysis since the measurements that are being compared are heterogeneous.

One well-studied example of the task of analysing paired heterogeneous biological data is gene expression quantitative trait loci (eQTL) mapping, which seeks to find associations between single nucleotide polymorphisms (SNPs) and transcript expression levels (Cookson *et al.*, 2009). A particularly interesting application of multivariate eQTL mapping relates to the discovery of association between biomarkers for ovarian cancer, which is treated as a case study in this work.

The application of paired analysis of heterogeneous biological data types is not limited to eQTL mapping, and involves the very common problem of discovering shared patterns of variation across different biological measurements. Other areas that have seen considerable interest include the association of genetic copy number variation with gene expression levels (Beroukhim *et al.*, 2007) and identification of genes for which DNA methylation is associated to gene expression (Louhimo and Hautaniemi, 2011). In these and related studies, the problem of comparing heterogeneous data types can be approached using a notion of distance between pairs of biological samples. Distances are scalar-valued measures of dissimilarity between two input observations, with greater values indicating greater dissimilarity between them. For a given application, a suitable semi-metric or metric distance measure can generally be defined depending on the nature of the data and on the specific objectives of the study. Such distances computed between all pairwise combinations of samples are arranged in square, symmetric matrices. These matrices then quantify the variability in the random sample according to that particular distance measure, and can be directly compared using an appropriate statistical test.

Distance-based approaches offer a number of advantages compared to direct comparisons involving the original measurements. Firstly, distances can overcome limitations of traditional multivariate approaches when dealing with high-dimensional random vectors observed on a much smaller number of sampling units (Shannon *et al.*, 2002). Secondly, distance measures can be appropriately defined for data generated under different experimental conditions, and are not confined to data types represented only by vectors. A number of distance measures are readily available for several types of biological data, including DNA sequences (Wu *et al.*, 1997), genetic markers such as SNPs (Goh *et al.*, 2011), gene expression data (Priness *et al.*, 2007), longitudinal gene expression data (Minas *et al.*, 2011) and proteins (Hollich *et al.*, 2005). Moreover, for a specific data type, more than one distance measure can be deployed to capture different aspects of the data. This can ultimately lead to the discovery of different types of associations.

Only a handful of statistical procedures are currently available to determine whether there is an association between paired distance matrices, and they all rely on computationally intensive procedures for statistical inference. The most commonly used and well-known procedure is the Mantel test (Mantel, 1967), a generalization of Pearson's correlation. The test has already been found useful in many applications arising in bioinformatics (Shannon *et al.*, 2002;

Beckmann *et al.*, 2005; Sun *et al.*, 2011). This testing procedure, however, has some limitations. First, when used to compare paired scalar-valued observations using the Euclidean distance, the test exhibits less power than a Pearson’s correlation test (Peres-Neto and Jackson, 2001; Legendre and Fortin, 2010). Secondly, computationally intensive Monte Carlo permutations are required to assess significance. This procedure, which assumes exchangeability of the sampling units under the null hypothesis of no association, introduces sampling errors which leads to inaccurate estimation of small p-values. This is especially true when too few permutations are used, generally between $O(10^3)$ and $O(10^5)$ (Berry and Mielke, 1983; Mielke and Berry, 2007; Knijnenburg *et al.*, 2009). For example, in order to obtain a permutation p-value within 10^{-5} of the true p-value, it has been shown that $O(10^7)$ permutations are required, which is computationally infeasible without access to specialised high-performance computational resources. This is the case in all situations where thousands or even millions of tests are required, such as in eQTL mapping studies in which the entire genome is scanned in search of potential associations. In such large-scale applications, it is not uncommon to settle for a much smaller number of Monte Carlo permutations per test, which in turn can lead to inflated family-wise type I error rates (Phipson and Smyth, 2010).

In this article we propose a novel test, the generalized RV (GRV) test, to detect association between paired distance matrices. This test is derived as a generalization of the classical multivariate RV test of no correlation between paired random vectors, originally proposed by Escoufier (1973). While inference can be performed via computationally intensive Monte Carlo permutations, we also derive an approximation to the exact null distribution which would be obtained by enumerating all permutations. Using the proposed null distribution, approximate p-values can be estimated in closed-form without the need for computationally intensive procedures. The proposed GRV test is introduced in Section 2. In Section 3 we describe power studies showing that the GRV test can offer greater statistical power than that achieved by the Mantel test even when many millions of Monte Carlo permutations are performed. Section 4 presents an application of the GRV test in cancer research, where the interest lies in detecting biological pathways characterized by a non-random association between genetic and gene expression sample variability. Concluding remarks are presented in Section 5.

2 The generalized RV test

2.1 Problem statement

Assume we have observed paired measurements $\{(x_i, y_i)\}_{i=1}^N$ on N independent biological sampling units. Each one of the $\{x_i\}_{i=1}^N$ and $\{y_i\}_{i=1}^N$ measurement sets is represented by a particular data structure, such as continuous or discrete vectors, although the data do not necessarily need to be represented as vectors. For instance, the measurements might be structured as graphs or trees (e.g. representing biological networks or ontologies). We also suppose that suitably chosen semi-metric or metric distances d_x and d_y have been identified that capture the dissimilarity between any sample pair $\{(x_i, x_j)\}_{i \neq j}$ and $\{(y_i, y_j)\}_{i \neq j}$, respectively. On evaluating the distances for all possible pairs, we obtain the paired $N \times N$ distance matrices $\Delta_x = \{d_x(x_i, x_j)\}_{i,j=1}^N$ and $\Delta_y = \{d_y(y_i, y_j)\}_{i,j=1}^N$.

Using the random sample, evidence of non-random association between these paired genomic measurements can be assessed by testing the null hypothesis that $\Delta_x \neq a\Delta_y$, for some positive constant a . This constant represents possible scaling differences between the

elements of each distance matrix, which may arise because of the chosen distance measures. Under this null hypothesis, the relationship between the elements in Δ_x is not maintained by the corresponding elements in Δ_y . The alternative hypothesis is that the distance matrices are equal up to a constant, in which case every distance $d_x(x_i, x_j)$ is linearly related to the paired distance $d_y(y_i, y_j)$. We are particularly concerned with settings where inferences must be drawn for a very large number of such tests simultaneously, hence the computational cost in obtaining a p-value for each test should be kept as low as possible.

2.2 The proposed test statistic

When the observations are vectorial, the mean-centered vectors $\{\mathbf{x}_i\}_{i=1}^N \in \mathbb{R}^P$ and $\{\mathbf{y}_i\}_{i=1}^N \in \mathbb{R}^Q$ can be arranged in paired data matrices, \mathbf{X} and \mathbf{Y} , of dimensions $N \times P$ and $N \times Q$, respectively. In this case, in order to establish whether the P and Q variables are correlated, the null hypothesis that $\Sigma_{xy} = \mathbf{0}$ is commonly tested, where Σ_{xy} is the $P \times Q$ covariance matrix containing the true covariances between the P variables observed in $\{\mathbf{x}_i\}_{i=1}^N$ and Q variables observed in $\{\mathbf{y}_i\}_{i=1}^N$. The RV statistic has been proposed for this task and arises as a generalization of Pearson's correlation (Escoufier, 1973). As with Pearson's correlation, RV is computed as the ratio of covariance to the square-rooted product of the variances,

$$\text{RV}(\mathbf{X}, \mathbf{Y}) = \frac{\text{tr}(\mathbf{X}^T \mathbf{Y} \mathbf{Y}^T \mathbf{X})}{\|\mathbf{X}^T \mathbf{X}\| \|\mathbf{Y}^T \mathbf{Y}\|}, \quad (1)$$

where $\text{tr}(\cdot)$ is the trace function and $\|\mathbf{A}\| = \sqrt{\text{tr}(\mathbf{A}^T \mathbf{A})}$ denotes the Frobenius norm for matrix \mathbf{A} . The RV statistic ranges between 0 and 1, with no association when $\mathbf{X}^T \mathbf{Y} = \mathbf{0}$, and perfect association when there exists a linear mapping which relates every P -dimensional observation in \mathbf{X} to every Q -dimensional observation in \mathbf{Y} . Thus larger values provide evidence against the null hypothesis. When $P = Q = 1$, $\text{RV}(\mathbf{X}, \mathbf{Y}) = \text{cor}^2(\mathbf{X}, \mathbf{Y})$, where $\text{cor}(\cdot, \cdot)$ denotes Pearson's correlation coefficient.

We first observe that the RV test can be interpreted in terms of the Euclidean distance matrices Δ_x and Δ_y arising by applying the Euclidean distance measure to $\{\mathbf{x}_i\}_{i=1}^N$ and $\{\mathbf{y}_i\}_{i=1}^N$, respectively. This is because, due to properties of the trace operator, $\text{tr}(\mathbf{X}^T \mathbf{Y} \mathbf{Y}^T \mathbf{X}) = \text{tr}(\mathbf{X} \mathbf{X}^T \mathbf{Y} \mathbf{Y}^T)$, $\|\mathbf{X}^T \mathbf{X}\| = \|\mathbf{X} \mathbf{X}^T\|$ and $\|\mathbf{Y}^T \mathbf{Y}\| = \|\mathbf{Y} \mathbf{Y}^T\|$, so that RV statistic (1) can be written in equivalent form as

$$\text{RV}(\mathbf{X}, \mathbf{Y}) = \frac{\text{tr}(\mathbf{X} \mathbf{X}^T \mathbf{Y} \mathbf{Y}^T)}{\|\mathbf{X} \mathbf{X}^T\| \|\mathbf{Y} \mathbf{Y}^T\|}. \quad (2)$$

Note that (1) differs from (2) in that emphasis is placed on the two symmetric $N \times N$ matrices $\mathbf{X} \mathbf{X}^T$ and $\mathbf{Y} \mathbf{Y}^T$, instead of the four covariance matrices $\mathbf{X}^T \mathbf{Y} \in \mathbb{R}^{P \times Q}$, $\mathbf{Y}^T \mathbf{X} \in \mathbb{R}^{Q \times P}$, $\mathbf{X}^T \mathbf{X} \in \mathbb{R}^{P \times P}$ and $\mathbf{Y}^T \mathbf{Y} \in \mathbb{R}^{Q \times Q}$. In settings where $P, Q > N$, which are common when comparing biological measurements, it is computationally convenient to use (2) over (1).

In addition, the $N \times N$ matrices $\mathbf{X} \mathbf{X}^T$ and $\mathbf{Y} \mathbf{Y}^T$ can each be written in terms of the Euclidean distance matrices Δ_x and Δ_y , respectively. This is due to the results that

$$\mathbf{X} \mathbf{X}^T = -\frac{1}{2} \mathbf{C} \Delta_x^2 \mathbf{C} \quad \text{and} \quad \mathbf{Y} \mathbf{Y}^T = -\frac{1}{2} \mathbf{C} \Delta_y^2 \mathbf{C},$$

where $\mathbf{C} = (\mathbf{I}_N - \mathbf{J}_N/N)$ is the symmetric $N \times N$ centering matrix where \mathbf{I}_N is the identity matrix of size N and \mathbf{J}_N is the square matrix of ones of size N (Gower, 1966; Borg and Groenen,

2005). Also, $-\frac{1}{2}\mathbf{C}\Delta_x^2\mathbf{C} = \mathbf{G}_x$, where \mathbf{G}_x is the square, symmetric and real-valued Gower's centered inner product matrix, and similarly for \mathbf{G}_y , so that $\text{tr}(\mathbf{X}\mathbf{X}^T\mathbf{Y}\mathbf{Y}^T) = \text{tr}(\mathbf{G}_x\mathbf{G}_y)$, $\|\mathbf{X}\mathbf{X}^T\| = \|\mathbf{G}_x\|$ and $\|\mathbf{Y}\mathbf{Y}^T\| = \|\mathbf{G}_y\|$. Thus RV is completely specified by the elements of the Euclidean distance matrices.

Based on these observations, we propose a generalization for the RV statistic that is entirely specified by any pair of distance matrices, not necessarily Euclidean, in order to test the null hypothesis that $\Delta_x \neq a\Delta_y$. The generalized RV (GRV) test statistic is defined as

$$\text{GRV}(\mathbf{G}_x, \mathbf{G}_y) = \frac{\text{tr}(\mathbf{G}_x\mathbf{G}_y)}{\|\mathbf{G}_x\| \|\mathbf{G}_y\|}, \quad (3)$$

noting the implicit assumption $\|\mathbf{G}_x\| \|\mathbf{G}_y\| > 0$ which is always satisfied in practice; $\|\mathbf{G}_x\| > 0$ since \mathbf{G}_x contains real-valued elements which are not all zero, and similarly for \mathbf{G}_y .

It is insightful to understand in which ways the proposed GRV statistic is similar to, and differs from, the classical Mantel statistic. The Mantel statistic is computed as the correlation between the upper-triangular distances of each distance matrix. Denote by \mathbf{v}_x the column vector containing the $A = N(N-1)/2$ upper-triangular distances $\{d_x(x_i, x_j)\}_{i>j}$, and similarly for \mathbf{v}_y . The Mantel statistic is then computed as $r_M(\Delta_x, \Delta_y) = \text{cor}(\mathbf{v}_x, \mathbf{v}_y)$, i.e.,

$$r_M(\Delta_x, \Delta_y) = \frac{1}{A-1} \sum_{i>j} \left(\frac{d_x(x_i, x_j) - \bar{x}}{s_x} \right) \left(\frac{d_y(y_i, y_j) - \bar{y}}{s_y} \right),$$

where $\bar{x} = \sum_{i>j} d_x(x_i, x_j)/A$, $s_x^2 = \sum_{i>j} (d_x(x_i, x_j) - \bar{x})^2/(A-1)$, and similarly for \bar{y} and s_y^2 . The statistic therefore standardizes the upper-triangular elements of each distance matrix to have zero mean and unit variance. In this way the distances between paired distance matrices can be directly compared.

It can be shown that, like Mantel, GRV is also a correlation coefficient. In particular, it can be proved that GRV is equal to the correlation between the vectorized matrices $\mathbf{G}_x/\|\mathbf{G}_x\|$ and $\mathbf{G}_y/\|\mathbf{G}_y\|$, i.e., the N^2 -dimensional vectors \mathbf{g}_x and \mathbf{g}_y with elements $\{g_x(x_i, x_j)/\|\mathbf{G}_x\|\}_{i,j=1}^N$ and $\{g_y(y_i, y_j)/\|\mathbf{G}_y\|\}_{i,j=1}^N$, respectively. That is, $\text{GRV}(\mathbf{G}_x, \mathbf{G}_y) = \text{cor}(\mathbf{g}_x, \mathbf{g}_y)$ (see Appendix A for a proof). While both Mantel and GRV are correlation coefficients, the difference between them lies in the methods of standardization applied to the distances in each case. In Mantel, the upper-triangular distances are subjected to a classical standardization, where their mean is subtracted and they are divided by their standard deviation. In GRV, however, all distance elements are considered, and they are squared, double-centered and normalized by dividing by their Frobenius norm. This difference leads to greater power of the GRV test to detect association between paired distance matrices than the Mantel test, as demonstrated in Section 3. Further properties of the GRV test statistic are provided in Appendix B, including a discussion of how it overcomes the limitation of Mantel for scalar-valued observations and Euclidean distances.

2.3 Approximate null distribution

Under the null hypothesis, the sampling distribution of GRV is generally unknown. This is because the quantity $T = \text{tr}(\mathbf{G}_x\mathbf{G}_y)$ in the numerator is completely specified by the elements of Δ_x and Δ_y , which have unknown distributions and are correlated. The null sampling distribution can be generated by using permutations of one of the centered inner product matrices, \mathbf{G}_y , say. For each of N_π permutations $\pi \in \Pi$, which are one-to-one mappings of the

Table 1: Power of the GRV and Mantel tests at a significance level of 0.1% for $N = \{30, 50, 70, 90, 100\}$ and all combinations of the genetic (gen) identity-by-state (IBS), Sokal and Sneath (SS), and Rogers and Tanimoto I (RTI) distances and gene expression (gex) Euclidean (Euc), Manhattan (Man) and Mahalanobis (Mah) distances. For the GRV test, the standard deviation of the power estimate is given in brackets. For the Mantel test, the confidence interval obtained with 95% coverage probability is given in square brackets. The average number of Monte Carlo permutations required to achieve the power estimates for Mantel are stated below the confidence interval in millions.

Gen dist	Gex dist	Test		N								
				30		50		70		90		100
IBS	Euc	GRV	0.126	(0.042)	0.351	(0.064)	0.566	(0.074)	0.712	(0.049)	0.780	(0.051)
		Mantel	0.081	[0.042, 0.126]	0.116	[0.056, 0.183]	0.225	[0.145, 0.291]	0.275	[0.228, 0.325]	0.322	[0.272, 0.374]
				3.081		1.575		9.921		13.51		19.61
	Man	GRV	0.139	(0.049)	0.375	(0.068)	0.590	(0.066)	0.730	(0.059)	0.772	(0.065)
		Mantel	0.069	[0.023, 0.120]	0.151	[0.091, 0.226]	0.240	[0.173, 0.306]	0.302	[0.238, 0.355]	0.378	[0.314, 0.444]
	Mah			1.859		2.134		6.290		13.80		8.794
		GRV	0.394	(0.061)	0.846	(0.047)	0.952	(0.030)	0.978	(0.020)	0.990	(0.014)
SS	Euc	Mantel	0.100	[0.049, 0.169]	0.311	[0.262, 0.355]	0.542	[0.513, 0.572]	0.715	[0.694, 0.735]	0.776	[0.761, 0.789]
				2.818		28.53		135.9		424.2		666.8
	Man	GRV	0.045	(0.028)	0.160	(0.046)	0.308	(0.068)	0.503	(0.073)	0.556	(0.068)
		Mantel	0.047	[0.018, 0.073]	0.121	[0.080, 0.171]	0.195	[0.133, 0.269]	0.310	[0.242, 0.387]	0.363	[0.290, 0.424]
	Mah			2.879		10.27		4.371		7.325		9.232
		GRV	0.092	(0.041)	0.481	(0.074)	0.784	(0.059)	0.907	(0.037)	0.927	(0.034)
RTI	Euc	Mantel	0.072	[0.030, 0.112]	0.266	[0.194, 0.340]	0.510	[0.448, 0.564]	0.652	[0.613, 0.687]	0.723	[0.689, 0.756]
				5.567		8.635		21.92		74.12		40.01
	Man	GRV	0.055	(0.035)	0.172	(0.052)	0.363	(0.067)	0.540	(0.076)	0.639	(0.072)
		Mantel	0.051	[0.019, 0.089]	0.124	[0.067, 0.171]	0.216	[0.154, 0.288]	0.360	[0.282, 0.432]	0.401	[0.321, 0.465]
	Mah			1.373		4.512		4.371		9.289		7.233
		GRV	0.173	(0.055)	0.675	(0.068)	0.925	(0.038)	0.987	(0.015)	0.993	(0.010)
		Mantel	0.062	[0.011, 0.119]	0.241	[0.180, 0.314]	0.515	[0.478, 0.554]	0.688	[0.672, 0.702]	0.765	[0.755, 0.776]
				0.8066		11.89		151.0		654.4		1068

set $\{1, \dots, N\}$ to itself, the rows and columns of \mathbf{G}_y are simultaneously permuted by π and denoted $\mathbf{G}_{y,\pi}$. This generates the set $\{\hat{\text{GRV}}(\mathbf{G}_x, \mathbf{G}_{y,\pi})\}_{\pi \in \Pi}$ and the p-value of an observed GRV statistic, $\hat{\text{GRV}}(\mathbf{G}_x, \mathbf{G}_y)$, can be obtained as the proportion of permuted statistic values greater than or equal to $\hat{\text{GRV}}(\mathbf{G}_x, \mathbf{G}_y)$. This is a right-tailed test as larger GRV values provide evidence against the null.

In order to estimate the p-value without permutations, we adopt a moment matching approach where the exact null distribution which would be obtained if all $N!$ permutations were used is approximated by a continuous distribution. In particular, we approximate the null distribution by the Pearson type III distribution, which has been shown to capture skewed characteristics of null sampling distributions (Mielke and Berry, 2007; Josse *et al.*, 2008). To use this distribution we require the first few moments of the exact permutation distribution of T ; the mean is $\mu = \frac{1}{N!} \sum_{\pi \in \Pi} \hat{T}_\pi$, the variance is $\sigma^2 = \frac{1}{N!} \sum_{\pi \in \Pi} \hat{T}_\pi^2 - \mu^2$ and the skewness is

$$\gamma = \frac{\frac{1}{N!} \sum_{\pi \in \Pi} \hat{T}_\pi^3 - 3\mu\sigma^2 - \mu^3}{\sigma^3},$$

where $\hat{T}_\pi = \text{tr}(\mathbf{G}_x \mathbf{G}_{y,\pi})$ and Π contains all $N!$ permutations. No permutations are needed to compute $\{\mu, \sigma, \gamma\}$, since closed-form expressions of these quantities are retrievable via the analytical results of Kazi-Aoual *et al.* (1995). These results require that \mathbf{G}_x and \mathbf{G}_y are square, symmetric and centered; properties satisfied by definition.

On obtaining closed-form expressions for the mean, variance and skewness of the exact permutation distribution of T , we standardize T by subtracting μ and dividing by σ . The resulting $T_s = (T - \mu)/\sigma$ is then assumed to have the Pearson type III distribution under the null with probability density function (PDF) denoted $f_{T_s}(t; \gamma)$; see Mielke and Berry (2007) for the full definition of the distribution. We denote the CDF of T_s by $\mathcal{F}_{T_s}(t; \gamma)$. The approximate null distribution of $\text{GRV}(\mathbf{G}_x, \mathbf{G}_y) = T/\|\mathbf{G}_x\| \|\mathbf{G}_y\|$ can then be derived by a simple transformation of the distribution of T_s . The CDF of GRV, $\mathcal{F}_{\text{GRV}}(x; \gamma)$, is given by

$$\mathcal{F}_{\text{GRV}}(x; \gamma) = \mathcal{F}_{T_s} \left(\frac{x \|\mathbf{G}_x\| \|\mathbf{G}_y\| - \mu}{\sigma}; \gamma \right),$$

and this is a valid CDF since $\mathcal{F}_{T_s}(\cdot; \gamma)$ is a valid CDF. The p-value of an observed GRV statistic $\hat{\text{GRV}}(\mathbf{G}_x, \mathbf{G}_y)$ is then estimated by $1 - \mathcal{F}_{\text{GRV}}(\hat{\text{GRV}}(\mathbf{G}_x, \mathbf{G}_y); \gamma)$. The PDF $f_{\text{GRV}}(x, \gamma)$ is found by differentiation as

$$f_{\text{GRV}}(x, \gamma) = \left(\frac{\|\mathbf{G}_x\| \|\mathbf{G}_y\|}{\sigma} \right) f_{T_s} \left(\frac{x \|\mathbf{G}_x\| \|\mathbf{G}_y\| - \mu}{\sigma}; \gamma \right),$$

since $\|\mathbf{G}_x\| \|\mathbf{G}_y\| > 0$ and $\sigma > 0$.

3 Power Studies

In this section we assess the power of the proposed GRV test and compare it to the classical Mantel test. A Monte Carlo procedure is used to estimate the power of the GRV test, and in order to estimate the power of the Mantel test we use the recently proposed algorithm of Gandy and Rubin-Delanchy (2013). This algorithm estimates the power of a Monte Carlo testing procedure and states a confidence interval around this estimate boasting a guaranteed coverage probability. It also states the average number of Monte Carlo permutations required

to achieve the given power estimate. It requires specification of the maximum length of the resulting confidence interval and the desired coverage probability, and then runs as many permutations as required to yield a power estimate with a confidence interval of length no greater than specified. Using this approach we are able to demonstrate that, at least for the experimental setups considered, many millions of permutations are needed for the Mantel test to achieve power of the same precision as the GRV test, and that the GRV test achieves greater power.

We use simulated data that mimics the experimental data of a typical eQTL application where transcriptional measurements are paired with SNPs. We first generate an $N \times P$ matrix \mathbf{X} containing N simulated SNPs (i.e., minor allele counts), denoted $\mathbf{x}_i = (x_{i1}, \dots, x_{iP})^T$ for $i = 1, \dots, N$, with varying minor allele frequencies (MAFs) at each of the P SNPs. The MAF of SNP p , m_p , is first simulated from a uniform distribution $U(0.1, 0.5)$, and the allele count is then simulated from a multinomial distribution with probabilities $(1 - m_p)^2$, $2m_p(1 - m_p)$, and m_p^2 of observing 0, 1 and 2 minor alleles, respectively. The paired data matrix $\mathbf{Y} = (\mathbf{y}_1, \dots, \mathbf{y}_N)^T$ is then generated as follows. The $N \times 1$ vector $\mathbf{z} = \mathbf{X}\mathbf{1}_P = (z_1, \dots, z_N)^T$ containing the row sums of \mathbf{X} , i.e., the minor allele counts across the P SNPs, is computed, and $\mathbf{y}_i = z_i\mathbf{1}_Q + \mathbf{e}_i$ for $i = 1, \dots, N$, where $\mathbf{e}_i \sim N_Q(\boldsymbol{\mu}, \boldsymbol{\Sigma})$, with $\boldsymbol{\mu} = (\mu_1, \dots, \mu_Q)^T$ and $\mu_q \sim U(0, 1)$ for $q = 1, \dots, Q$, and $\boldsymbol{\Sigma}$ a random $Q \times Q$ Wishart matrix. We then consider the widely-used identity-by-state (IBS) distance for the simulated genetic markers in \mathbf{X} , which measures the degree to which alleles are shared across the SNPs (see, for instance, Goh *et al.* (2011)), in addition to the Sokal and Sneath, and Rogers and Tanimoto I distances (Selinski and Ickstadt, 2005). For the observations in \mathbf{Y} we use the Euclidean, Manhattan and Mahalanobis distances. See Appendix C for details on these distances.

For $P = 2$, $Q = 10$, and each of $N = \{30, 50, 70, 90, 100\}$, the power of the GRV test is computed using 50 Monte Carlo runs, and each time generating 50 paired datasets as above. For each genetic and gene expression distance, the GRV test is applied and the proportion of p-values less than or equal to the significance level of 0.1% is recorded. The mean power across all 50 Monte Carlo runs for each distance combination and value of N is reported in Table 1, in addition to the standard deviation of these estimates. As expected, the power increases with N .

We then estimate the power of the Mantel test for each of the above settings using the algorithm of Gandy and Rubin-Delanchy (2013). For each setting we bound the confidence interval length by twice the standard deviation of the corresponding estimate for the GRV test. That is, we consider one standard deviation on either side of the power estimate as an empirical indication of the precision achieved by the GRV test. We monitor the number of Monte Carlo permutations required to obtain power estimates with such confidence intervals with a coverage probability of 95%. These results are also given in Table 1, where the power is stated alongside the confidence interval, and the number of Monte Carlo permutations required is stated on a separate line below the confidence interval.

In all settings the GRV test achieves greater power than the Mantel test, even when using large numbers of Monte Carlo permutations. Note that while the Mantel power estimates improve with N , as expected, the number of Monte Carlo permutations varies between $O(10^5)$ to $O(10^9)$ non-linearly with N . This is an artefact of the algorithm; the expected number of permutations depends on the length of the confidence interval sought and the region of the confidence interval in $[0, 1]$ (see Gandy and Rubin-Delanchy (2013) for more details). An example of how the standardization used by GRV yields greater power than that used by

Mantel is presented in Appendix D, in addition to simulation results demonstrating that the GRV test attains the nominal significance level of a given test.

4 Application to ovarian cancer

Around 225,000 women were estimated to have developed ovarian cancer in 2008, of which over 80% are Epithelial Ovarian Cancer (EOC). In the UK EOC has particularly poor prognosis with only 34% of patients surviving 5 years after diagnosis.

In many cancers, the ability to measure some biomarker that can predict the likely response to a chemotherapeutic agent could prevent patients unnecessarily suffering side-effects from ineffective treatments, and may suggest the best course of action where multiple treatments are available. Unfortunately, while high-throughput gene expression profiling has been successful in identifying prognostic ‘signatures’ for many cancers (van’t Veer and Bernards, 2008), solid tumour gene expression profiles have practical drawbacks preventing their use as prognostic or predictive biomarkers in the clinic, especially in terms of obtaining mRNA from the tumour and the transitive nature of gene expression (Sawyers, 2008). However, circulating tumour DNA can be isolated from serum samples in ovarian cancer patients (Hickey, 1999), which reflect an attractive source for genetic biomarkers both because of the relative ease of accessibility of the patient material and the long-term stability of the genetic sequence. Additionally, treating biological pathways as multidimensional biomarkers may provide substantial improvements in terms of robustness over single agents (Dai *et al.*, 2011).

We sought to discover if it might be possible to use genetic biomarkers as surrogates for the transcriptional activation of pathways. Such pathway-based genetic biomarkers may suggest starting points for the development of clinical diagnostic tests to predict the outcome of a therapeutic exploiting a related pathway.

In order to identify pathways for which the profile of SNP genotypes across genes in the pathway closely reflects the state of expression of the genes in the pathway, we applied the GRV test to evaluate the significance of the degree of similarity in distance matrices derived from SNP genotypes and gene expression levels from a large set of selected pathways, for two independent cohorts.

4.1 Datasets

For this study we use two independent cohorts of ovarian cancer patients: one collected by the TCGA consortium, and an independent cohort for validation. For the TCGA cohort, raw data CEL files from Affymetrix SNP 6.0 and from Affymetrix HT-HGU133A microarray profiling 499 primary high grade serous ovarian tumours were obtained from the TCGA web repository. For the Affymetrix SNP 6.0 arrays, normalization and genotype calling were performed using the CRLMM method (Carvalho *et al.*, 2007). Current annotations were obtained from Affymetrix’s NetAffx Analysis Center and SNPs were annotated with their associated Gene Symbols. Once SNP probes were mapped to Gene Symbols, lists of SNPs corresponding to each CPDB pathway were obtained by selecting all SNP probes mapping to Gene Symbols that were included in the pathway’s list of `hgnc_symbol_ids` as downloaded from Consensus-PathDB (CPDB) (Kamburov *et al.*, 2011). For the Affymetrix HT-HGU133A arrays, data normalization was performed using the RMA method (Irizarry *et al.*, 2003). Probe-set mapping to Gene Symbols was performed using the `hthgu133a.db` package from Bioconductor. Once probe-sets were mapped to Gene Symbols, lists of probe-sets corresponding to each

CPDB pathway were obtained by selecting all probe-sets mapping to Gene Symbols that were included in the pathway's list of `hgnc_symbol_ids` as downloaded from CPDB.

As an independent validation data set, raw data CEL files from both Affymetrix SNP 6.0 and Affymetrix HGU133plus2 microarrays profiling 60 primary high grade serous ovarian tumours were obtained from the Hammersmith Biobank Resource Centre. The datasets were provided courtesy of the Genome Institute of Singapore. Raw data for this independent tumour cohort was processed identically to those obtained from TCGA, although mapping of gene expression probe-sets to Gene Symbols was performed using the `hgu133plus2.db` package from Bioconductor. The tables of genotype calls and normalized gene expression values, mapped to pathways which were used in this analysis, are provided in Supplementary Tables S1 and S2, respectively.

We wanted to reduce the possibility that any identified associations simply reflected genome-wide similarity structures between the patients (e.g. due to molecular subtypes of disease) or involved so few measurements that they weren't indicative of the entire pathway. From the full list of 4,119 CPDB pathways we filtered out all those with less than 5 or more than 200 probe sets in either gene expression or SNP genotype array datasets. A total of 479 pathways remained after this filtering, and these pathways are listed in Supplementary Table S3 along with their mappings to probes in each of the two datasets.

4.2 Experimental results

Since different distance measures highlight different relationships among genetic markers and features of gene expression data (Priness *et al.*, 2007; Song *et al.*, 2012), we sought to perform a meta-analysis to combine the results of similarity testing across a range of distance measures for both the genotypes and gene expression values. The rationale being that if a given pathway shows similar distance matrices at the genotype and gene expression level for a number of different distance measures, the observed similarity is unlikely to be due to a quirk of one particular distance measure. For each of the 479 pathways, sample-wise distance matrices were calculated using 8 different distance measures (Euclidean, Mahalanobis, Manhattan, Maximum, Bray-Curtis, Pearson's correlation, Spearman's correlation and Cosine angle distances) for the gene expression data and 5 distance measures (IBS, SS, RTI, Simple Matching and Hamman I) for the SNP genotypes; see Appendix C for details on these distances. The association between the paired distance matrices, one derived from the genotypes of SNPs mapping to genes in the pathway and the other derived from the expression levels of genes in the pathway, was evaluated using the GRV test. This resulted in 40 individual p-values for each pathway, one for each combination of distance measures used. An illustration of how the null distribution compares with the permutation distribution when applied to this cancer data is presented in Appendix E. The 40 individual p-values for each pathway were then combined using the 'maximum' method as implemented in the R package *MetaDE*. The combined significance estimates from running this meta-analysis for each pathway are given in Supplementary Table S4. This meta-analysis was then repeated for the validation cohort.

To obtain a perspective of the robustness/repeatability of the associations detected between genotype and gene expression at the pathway level across the two independent cohorts, we compared the ranked lists of pathways arising from the meta-analysis of each cohort. The degree of agreement between the two ranked lists was evaluated using the normalized Canberra distance (Jurman *et al.*, 2008). The general principle is to quantify the distance between the rankings of the top k terms of two ranked lists, with each term's influence on the

overall distance weighted by its position in its respective list. By evaluating the normalized Canberra distance for the top-ranked k pathways, the relationship between overlap and ranking can be explored: smaller values indicate a lesser discrepancy in rankings among the top k terms and therefore more similar rankings in the two lists. Figure 1 shows the normalized Canberra distance plotted against k .

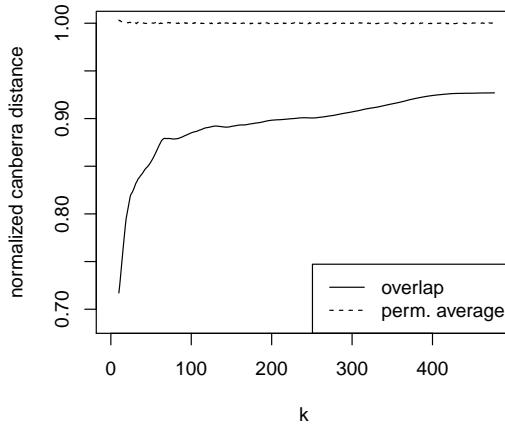


Figure 1: Plot of the normalized Canberra distance against k for the top k elements in two lists of pathways. Each pathway list is ranked according to meta-analysis estimate of significance of association between genotype and gene expression over 40 combinations of distance measures, one ranked list per cohort. The dotted line represents the average normalized Canberra distance between the TCGA data-derived pathway list and each of 5,000 random permutations of the ranks in the second list.

It can be seen that the overlap of the meta-analysis results from the two cohorts is greatest for small k , implying that the highest-ranking pathways in each list are more similar than the lower-ranking pathways. Furthermore, for each k we used 5,000 permutations of the ranks in the second list to obtain a p-value estimate for the normalized Canberra distance values observed between the two rankings. These p-values across the range of k were less than 0.005 after false discovery rate multiple-testing adjustment, reflecting the fact that the overlap between the two ranked lists of pathways was significantly greater than that expected by chance. These results suggest that the association between distance matrices based on the levels of expression of genes within certain pathways and distance matrices based on SNP genotypes associated with genes in the corresponding pathways is robust and reliably evaluated by the GRV test. The meta-analysis of the GRV test result across multiple distance measures demonstrates that these associations are not a result of peculiarities of particular distance measures, and certain pathways are shown to exhibit a very strong association in both cohorts.

A particularly striking result of this study is the CPDB pathway ‘5-aminoimidazole ribonucleotide biosynthesis’, which is ranked #1 in the meta-analysis of GRV test results from the TCGA dataset and ranked #2 in the meta-analysis of GRV test results from the validation dataset. The 5-aminoimidazole ribonucleotide is a precursor to de novo purine biosynthesis, which is essential for rapidly proliferating tissue (Firestone and Davisson, 1993). Pharmacological inhibition of de novo purine biosynthesis has proved to be a successful strategy in the

clinic for treating cancers and inflammatory disorders (Christopherson *et al.*, 2002). It has been shown through a clinical trial that genetic polymorphisms could be biomarkers to predict response to methotrexate therapy, which inhibits de novo purine biosynthesis, in patients with rheumatoid arthritis (Dervieux *et al.*, 2004). The ‘methotrexate pathway’ itself showed significant genotype-expression association in the TCGA cohort (ranked #17), but not in the validation cohort (ranked #316). The ‘5-aminoimidazole ribonucleotide biosynthesis’ pathway may therefore suggest additional or alternative SNPs to use as genetic biomarkers for predicting outcome of inhibiting de novo purine biosynthesis.

Another result of note is that the CPDB pathway Platinum Pathway, Pharmacokinetics/Pharmacodynamics yields highly significantly similar tumour-wise distance matrices across all 40 applications of the GRV test on this paired ovarian cancer data, for both cohorts investigated (ranked #14 from the TCGA dataset and #42 from the validation dataset). This is particularly interesting because platinum-based chemotherapy is the standard treatment for ovarian cancer, but unlike for many other solid cancers, survival rates in ovarian cancer have not improved noticeably in the 30 years since this treatment was introduced (Vaughan *et al.*, 2011). The principal reason for this is the high frequency of platinum-resistant relapse, which may occur through multiple mechanisms (Agarwal and Kaye, 2003; Stronach *et al.*, 2011). Therefore, being able to predict the response to the platinum-based DNA damaging agents could be of great clinical value in the treatment of ovarian cancer. Additionally, all genotype-derived and expression-derived distance matrices for the Base Excision Repair pathway are highly significantly similar in the TCGA cohort. This is of interest as the Base Excision Repair pathway is involved in repairing platinum-induced DNA damage. In fact, effective inhibition of this pathway (via PARP) in patients with BRCA1/2 mutations is associated with sensitivity to platinum chemotherapy (Fong *et al.*, 2010).

5 Conclusion

In this work we have presented the GRV test as a novel procedure to detect association between paired distance matrices, which is applicable in any setting where suitable distance measures can be defined. Similarly to the widely used Mantel test, the GRV test can be directly applied whether the distance measures are metric or semi-metric, but overcomes the limitation of the Mantel test that, where paired data is vector-valued, a hypothesis of no correlation between the vectors of interest can be tested. As a further advantage of the GRV test over existing distance-based tests, an approximate null distribution is proposed such that inference can be performed without expensive Monte Carlo permutation procedures. Through extensive simulations we have demonstrated that the GRV test achieves greater power than the popular Mantel test. This greater power is exhibited even when many millions of permutations are used for the Mantel test, thus demonstrating its suitability as a tool for use in multiple-testing settings where many tests need to be simultaneously performed.

The GRV test was applied to paired ovarian cancer datasets in which genome-wide profiles of SNP genotypes and gene expression levels were available for each patient across two independent cohorts ($N = 499$ and $N = 60$). The paired datasets were split into subsets in which the features all mapped to the same pathway (using mappings from the Consensus Pathway DB), so that each GRV test result would indicate the statistical significance of the similarity between the (tumour-wise) distance matrix derived from the genotype data and the distance matrix derived from the gene expression data. For each cohort a meta-analysis was

performed across the GRV test results from 40 different combinations of distance measures, and the pathways were ranked according to their significance across all distance measures. Permutation-based estimates of the significance of the overlap between the ranked lists of pathways from the two independent cohorts revealed a highly significant agreement in the two rankings, indicating a degree of robustness in the results of our analysis.

This application of the GRV test has demonstrated that the transcriptional activity of a number of pathways seem to be reliably predicted by sequencing a limited set of genetic markers that could be detected in circulating tumour DNA obtained from patients' serum. This relationship was particularly pronounced for the pathway '5-aminoimidazole ribonucleotide biosynthesis', for which the genotypes and the gene expression patterns were strikingly associated in two independent cohorts of patients. The association between genotypes and pathway-level gene expression profile is of particular interest in the case of the platinum response pathway, which is clearly clinically important in EOC given that the only standard treatment involves platinum-based chemotherapy. Following further investigation, these observations could potentially lead to development of non-invasive biomarkers that could predict patient response to front-line chemotherapy, and serve to stratify patients for selection into clinical trials involving alternative treatments.

As individual distance measures can yield different results in the analysis of associations across gene expression patterns, unless a single distance measure is selected *a priori* as the most appropriate for each dataset, combining results of multiple tests of association using different distance measures may provide more reliable results than a single test. A key advantage of the GRV test in our application to multivariate eQTL mapping of pathways in ovarian cancer is that it enables a meta-analysis of the associations between genetic distance and gene expression distance for different combinations of distance measures. Having a test for these associations that does not require expensive permutations enables a fast estimation of the robustness of the genotype-expression pathway level associations against quirks of particular definitions of distance. Given the robust associations discovered in this analysis, we suggest that it may be possible to predict the mode of activation of a pathway (in terms of which members are transcriptionally activated) in cancer patients, based on the genotype of the tumours' DNA across a limited panel of SNPs.

Acknowledgements

The authors would like to thank Professor Edison T. Liu, Genome Institute of Singapore, and the Agency for Science Technology and Research of Singapore for use of ovarian cancer microarray data. E. C. was funded by Ovarian Cancer Action, and G. M. and C. M. acknowledge financial support from the EPSRC.

A GRV statistic is a correlation coefficient

Consider the quantity $\text{cor}(\mathbf{g}_x, \mathbf{g}_y)$. To compute this we require the mean and standard deviation of the values in \mathbf{g}_x and \mathbf{g}_y . The means are 0 since \mathbf{G}_x and \mathbf{G}_y are centered matrices.

The standard deviation of the elements in \mathbf{g}_x is given by

$$\begin{aligned}
\sqrt{\frac{1}{N^2-1} \sum_{i=1}^N \sum_{j=1}^N \left(\frac{g_x(x_i, x_j)}{\|\mathbf{G}_x\|} \right)^2} &= \sqrt{\frac{1}{\|\mathbf{G}_x\|^2 (N^2-1)} \sum_{i=1}^N \sum_{j=1}^N g_x^2(x_i, x_j)} \\
&= \sqrt{\frac{\|\mathbf{G}_x\|^2}{\|\mathbf{G}_x\|^2 (N^2-1)}} \\
&= \sqrt{\frac{1}{N^2-1}},
\end{aligned}$$

and similarly for \mathbf{g}_y . Thus the correlation of interest is given by

$$\begin{aligned}
\text{cor}(\mathbf{g}_x, \mathbf{g}_y) &= \frac{1}{N^2-1} \sum_{i=1}^N \sum_{j=1}^N \left(\frac{g_x(x_i, x_j)/\|\mathbf{G}_x\|}{\sqrt{\frac{1}{N^2-1}}} \right) \left(\frac{g_y(y_i, y_j)/\|\mathbf{G}_y\|}{\sqrt{\frac{1}{N^2-1}}} \right) \\
&= \frac{1}{N^2-1} \frac{1}{\left(\frac{1}{N^2-1}\right)} \sum_{i=1}^N \sum_{j=1}^N \frac{g_x(x_i, x_j)g_y(y_i, y_j)}{\|\mathbf{G}_x\| \|\mathbf{G}_y\|} \\
&= \frac{1}{\|\mathbf{G}_x\| \|\mathbf{G}_y\|} \sum_{i=1}^N \sum_{j=1}^N g_x(x_i, x_j)g_y(y_i, y_j) \\
&= \frac{\text{tr}(\mathbf{G}_x \mathbf{G}_y)}{\|\mathbf{G}_x\| \|\mathbf{G}_y\|} \\
&= \text{GRV}(\mathbf{G}_x, \mathbf{G}_y),
\end{aligned}$$

as required.

B Properties of the GRV statistic

The interpretation of GRV as a correlation coefficient indicates that it may range between -1 and 1 , with negative values indicating association in the form of a linear correlation of different sign (as with Pearson's correlation and Mantel). In this section we discuss how to interpret the values of GRV.

To begin, we show that the GRV statistic is related to the Frobenius distance between \mathbf{G}_x and \mathbf{G}_y . Recall that the Frobenius distance between two matrices \mathbf{A} and \mathbf{B} of equal dimensions is defined by $d_F(\mathbf{A}, \mathbf{B}) = \sqrt{\text{tr}((\mathbf{A} - \mathbf{B})^T(\mathbf{A} - \mathbf{B}))}$. It is easily shown that the Frobenius distance between the scale invariant configurations $\mathbf{G}_x/\|\mathbf{G}_x\|$ and $\mathbf{G}_y/\|\mathbf{G}_y\|$ is related to the GRV statistic by

$$d_F \left(\frac{\mathbf{G}_x}{\|\mathbf{G}_x\|}, \frac{\mathbf{G}_y}{\|\mathbf{G}_y\|} \right) = \sqrt{2(1 - \text{GRV}(\mathbf{G}_x, \mathbf{G}_y))}. \quad (4)$$

When $d_F(\mathbf{G}_x/\|\mathbf{G}_x\|, \mathbf{G}_y/\|\mathbf{G}_y\|) = 0$, observe that perfect association is achieved when $\text{GRV}(\mathbf{G}_x, \mathbf{G}_y) = 1$, i.e., when

$$\frac{\mathbf{G}_x}{\|\mathbf{G}_x\|} = \frac{\mathbf{G}_y}{\|\mathbf{G}_y\|}. \quad (5)$$

This equality, however, can only be attained if \mathbf{G}_x and \mathbf{G}_y are both positive semi-definite (having non-negative diagonals), or both indefinite (having non-negative and negative values on the diagonals). These occur if d_x and d_y are both metric, or semi-metric, respectively. When one distance function is metric and the other is semi-metric, perfect association cannot be attained, as the diagonals of \mathbf{G}_x and \mathbf{G}_y cannot be equal. In addition, negative values are attained for GRV for certain combinations of these metric properties.

To see this, consider the upper and lower bounds of the GRV statistic, found as follows. Let the ordered eigenvalues of \mathbf{G}_x and \mathbf{G}_y be denoted $\{\lambda_{x,i}\}_{i=1}^N$ and $\{\lambda_{y,i}\}_{i=1}^N$, respectively, and consider the bounds of the quantity $\text{tr}(\mathbf{G}_x \mathbf{G}_y)$. Using the result of Lasserre (1995), which gives bounds for the trace of the product of two Hermitian matrices (square, complex-valued, and equal to their conjugate transpose), these are given by

$$\sum_{i=1}^N \lambda_{x,i} \lambda_{y,N-i+1} \leq \text{tr}(\mathbf{G}_x \mathbf{G}_y) \leq \sum_{i=1}^N \lambda_{x,i} \lambda_{y,i}.$$

The bounds for GRV are then given by

$$\frac{\sum_{i=1}^N \lambda_{x,i} \lambda_{y,N-i+1}}{\|\mathbf{G}_x\| \|\mathbf{G}_y\|} \leq \text{GRV}(\mathbf{G}_x, \mathbf{G}_y) \leq \frac{\sum_{i=1}^N \lambda_{x,i} \lambda_{y,i}}{\|\mathbf{G}_x\| \|\mathbf{G}_y\|},$$

since $\|\mathbf{G}_x\| \|\mathbf{G}_y\| > 0$. The numerators of the bounds of GRV are sums of eigenvalue products, and so may be non-negative or negative depending on the sign of the eigenvalues. This in turn depends on the distance functions satisfying the metric property; they are either both metric or both semi-metric, or one is metric and the other is semi-metric. Note that greater Frobenius distance values occur with lower GRV values which may be negative, and so the paired distance matrices are considered to be less associated.

B.1 Both metric or semi-metric distances

When the distances are both metric, \mathbf{G}_x and \mathbf{G}_y are positive semi-definite and the ordered eigenvalues $\{\lambda_{x,i}\}_{i=1}^N$ and $\{\lambda_{y,i}\}_{i=1}^N$ are non-negative (Krzanowski, 2000). The summation in the lower bound contains the terms $\{\lambda_{x,i} \lambda_{y,N-i+1}\}_{i=1}^N$, which are therefore non-negative, so that GRV is non-negative. The minimum value of 0 indicating no association is attained when $\text{tr}(\mathbf{G}_x \mathbf{G}_y) = 0$. This can be interpreted in terms of the $N \times N$ principal coordinate matrices arising from Δ_x and Δ_y , denoted $\tilde{\mathbf{X}}$ and $\tilde{\mathbf{Y}}$, respectively. These matrices are found by classical multidimensional scaling (MDS) such that the rows are N -dimensional representations of the original observations in Euclidean space, and their pairwise Euclidean distances equal the corresponding pairwise distances in Δ_x and Δ_y (Torgerson, 1952; Gower, 1966). It then follows that $\mathbf{G}_x = \tilde{\mathbf{X}} \tilde{\mathbf{X}}^T$ and $\mathbf{G}_y = \tilde{\mathbf{Y}} \tilde{\mathbf{Y}}^T$, so that $\text{tr}(\mathbf{G}_x \mathbf{G}_y) = \text{tr}(\tilde{\mathbf{X}} \tilde{\mathbf{X}}^T \tilde{\mathbf{Y}} \tilde{\mathbf{Y}}^T) = 0 \Rightarrow \tilde{\mathbf{X}}^T \tilde{\mathbf{Y}} = \mathbf{0}$, i.e., the principal coordinate matrices are orthogonal. The maximum value GRV can take is 1, since \mathbf{G}_x and \mathbf{G}_y have positive diagonal elements so that equality (5) can be attained. In this case the distance matrices are equal up to a positive scaling factor, and there is perfect association.

Note also that when \mathbf{X} and \mathbf{Y} are centered vector-valued observations with $\mathbf{X}^T \mathbf{Y} = \mathbf{0}$, and d_x and d_y are the Euclidean distance functions, GRV yields a value of 0. This is because $\tilde{\mathbf{X}} \tilde{\mathbf{X}}^T = \mathbf{X} \mathbf{X}^T$ and $\tilde{\mathbf{Y}} \tilde{\mathbf{Y}}^T = \mathbf{Y} \mathbf{Y}^T$, so that $\text{tr}(\mathbf{G}_x \mathbf{G}_y) = \text{tr}(\tilde{\mathbf{X}} \tilde{\mathbf{X}}^T \tilde{\mathbf{Y}} \tilde{\mathbf{Y}}^T) = \text{tr}(\mathbf{X} \mathbf{X}^T \mathbf{Y} \mathbf{Y}^T)$. This result applies when the data is scalar-valued, so GRV equals 0 inline with the raw data being uncorrelated. Thus GRV overcomes the limitation of Mantel.

On the other hand, when the distances functions are both semi-metric, the ordered eigenvalues $\{\lambda_{x,i}\}_{i=1}^N$ and $\{\lambda_{y,i}\}_{i=1}^N$ are both non-negative and negative, so that GRV may attain negative values. Since the diagonal elements of \mathbf{G}_x and \mathbf{G}_y are both positive and negative, there may exist two such matrices with equal diagonals. Hence there may exist a scenario in which equality (5) is attained, and perfect association is indicated by a GRV value of 1.

B.2 Metric and semi-metric distances

When one of the distances, d_x for instance, is metric while d_y is semi-metric, then only the ordered eigenvalues $\{\lambda_{x,i}\}_{i=1}^N$ are strictly non-negative, so that the summation in the lower bound may be negative. In this case, GRV may attain negative values. The maximum attainable value of GRV is not 1 in this case, since equality (5) cannot be attained. This is because the diagonals of \mathbf{G}_x are positive while the diagonals of \mathbf{G}_y are both positive and negative. Thus perfect association cannot be attained, but larger values indicate greater association.

This result is not surprising because the relationship between pairwise distances with respect to d_x and d_y cannot be the same. Recall that metric distance functions satisfy the triangle inequality so that $d_x(x_i, x_j) \leq d_x(x_i, x_k) + d_x(x_k, x_j)$ for any three observations x_i, x_j, x_k (Mardia *et al.*, 1979). The corresponding distances with respect to d_y will not share this property as d_y is semi-metric, that is, $d_y(y_i, y_j)$ does not satisfy $d_y(y_i, y_j) \leq d_y(y_i, y_k) + d_y(y_k, y_j)$ for y_i, y_j, y_k . Thus the inter-point relationships between all the distances in Δ_x will not match those in Δ_y (for if they did d_y would satisfy the triangle inequality and hence be metric).

C Distance measures

C.1 SNPs

Assume two P -dimensional vectors $\mathbf{x} = (x_1, \dots, x_P)^T$ and $\mathbf{y} = (y_1, \dots, y_P)^T$ with discrete-valued elements representing minor allele counts at P SNPs. The identity-by-state (IBS) distance measure is defined as

$$d_{IBS}(\mathbf{x}, \mathbf{y}) = 1 - \frac{1}{2P} \sum_{p=1}^P s(x_p, y_p),$$

where $s(x_p, y_p) = 0$ if $x_p = 0$ and $y_p = 2$, or if $x_p = 2$ and $y_p = 0$, $s(x_p, y_p) = 1$ if $x_p = 1$ and $y_p \neq 1$, or if $y_p = 1$ and $x_p \neq 1$, and $s(x_p, y_p) = 2$ if $x_p = y_p$. This distance takes values between 0 and 1 and is semi-metric.

Genetic distances have also been proposed based on the contingency table containing the frequency that each combination of minor allele counts occurs over the SNPs (Selinski and Ickstadt, 2005); see Table 2.

The key statistics in this table are the number of complete matches of the minor allele counts, $m^+ = \sum_{k=0}^2 m_{kk}$, and the number of mismatches, $m^- = P - m^+$, where the total number of possible matches is P . Based on these quantities, the following ‘matching coefficient’ distance measures can be defined; the Simple Matching distance

$$d_{SM}(\mathbf{x}, \mathbf{y}) = 1 - \frac{m^+}{P},$$

Table 2: Contingency table containing the frequency of a given combination of minor allele count between \mathbf{x} and \mathbf{y} over the P SNPs. m_{kl} is the frequency of \mathbf{x} having k minor alleles and \mathbf{y} having l minor alleles.

		\mathbf{y}	0	1	2
			0	1	2
\mathbf{x}	0	m_{00}	m_{01}	m_{02}	
	1	m_{10}	m_{11}	m_{12}	
	2	m_{20}	m_{21}	m_{22}	

the Sokal and Sneath (SS) distance

$$d_{SS}(\mathbf{x}, \mathbf{y}) = 1 - \frac{m^+}{m^+ + \frac{1}{2}m^-},$$

and the Rogers and Tanimoto I (RTI) distance

$$d_{RTI}(\mathbf{x}, \mathbf{y}) = 1 - \frac{m^+}{m^+ + 2m^-}.$$

There is also the Hamman I similarity measure

$$s_{HI}(\mathbf{x}, \mathbf{y}) = \frac{m^+ - m^-}{P},$$

which can be transformed into a distance measure as follows. Assume N P -dimensional minor allele count vectors $\{\mathbf{x}_i\}_{i=1}^N$; this is required in order to normalize the magnitude of the similarities to the range $[0, 1]$. The Hamman I distance between \mathbf{x}_i and \mathbf{x}_j is then given by

$$d_{HI}(\mathbf{x}_i, \mathbf{x}_j) = 1 - \frac{s^*(\mathbf{x}_i, \mathbf{x}_j)}{\max_{i,j}\{s^*(\mathbf{x}_i, \mathbf{x}_j)\}},$$

where $s^*(\mathbf{x}_i, \mathbf{x}_j) = s_{HI}(\mathbf{x}_i, \mathbf{x}_j) + |\min_{i,j}\{s_{HI}(\mathbf{x}_i, \mathbf{x}_j)\}|$. This takes values between 0 and 1 and is semi-metric.

C.2 Vectors

Assume two P -dimensional real-valued vectors $\mathbf{x} = (x_1, \dots, x_P)^T$ and $\mathbf{y} = (y_1, \dots, y_P)^T$. Many measures exist (see, for example, Pekalska and Duin (2005)), of which a few are provided in Table 3, along with their ranges and properties, i.e., whether they are metric or semi-metric. These include the Euclidean, Manhattan, Maximum, Bray-Curtis, Pearson's correlation and the Cosine angle distances.

Spearman's correlation distance is defined by applying Pearson's correlation to the ranks of the elements of the vectors, rather than the actual values. In particular, let $\mathbf{x}_r = (x_{r1}, \dots, x_{rP})^T$ and $\mathbf{y}_r = (y_{r1}, \dots, y_{rP})^T$ be the vectors containing the ranks of the elements of \mathbf{x} and \mathbf{y} , respectively in ascending order (highest value given rank 1). That is, x_{rp} is the rank of x_p , and similarly for y_{rp} . If several elements of a given vector are equal, they are

Table 3: Commonly encountered distance measures for vector-valued objects. The (M) or (SM) by each distance name indicates whether it is metric or semi-metric.

Distance	Notation	Definition	Range
Euclidean (M)	$d_E(\mathbf{x}, \mathbf{y})$	$\sqrt{(\mathbf{x} - \mathbf{y})^T (\mathbf{x} - \mathbf{y})}$	$[0, \infty)$
Manhattan (SM)	$d_{MAN}(\mathbf{x}, \mathbf{y})$	$\sum_{p=1}^P x_p - y_p $	$[0, \infty)$
Maximum (SM)	$d_{MAX}(\mathbf{x}, \mathbf{y})$	$\max_p \{ x_p - y_p \}$	$[0, \infty)$
Bray-Curtis (SM)	$d_{BC}(\mathbf{x}, \mathbf{y})$	$\frac{\sum_{p=1}^P x_p - y_p }{\sum_{p=1}^P (x_p + y_p)}$	$[0, \infty)$
Mahalanobis (SM)	$d_{MAH}(\mathbf{x}, \mathbf{y})$	$\sqrt{(\mathbf{x} - \mathbf{y})^T \mathbf{S}^{-1} (\mathbf{x} - \mathbf{y})},$ \mathbf{S} a $P \times P$ covariance matrix, $P < N$	$[0, \infty)$
Pearson's correlation (SM)	$d_{PC}(\mathbf{x}, \mathbf{y})$	$1 - \frac{\sum_{p=1}^P (x_p - \bar{x})(y_p - \bar{y})}{\sum_{p=1}^P (x_p - \bar{x})^2 \sum_{p=1}^P (y_p - \bar{y})^2},$ $\bar{x} = \frac{1}{P} \sum_{p=1}^P x_p, \quad \bar{y} = \frac{1}{P} \sum_{p=1}^P y_p$	$[0, 2]$
Cosine angle (SM)	$d_{PC}(\mathbf{x}, \mathbf{y})$	$1 - \frac{\mathbf{x}^T \mathbf{y}}{\ \mathbf{x}\ \ \mathbf{y}\ },$ $\ \mathbf{x}\ = \sqrt{\sum_{p=1}^P x_p^2}, \quad \ \mathbf{y}\ = \sqrt{\sum_{p=1}^P y_p^2}$	$[0, 2]$

assigned a rank equal to the mean of their respective positions in the list of ascending values. For example, for vector $(0.1, 0.4, 0.4, 0.5, -31)^T$, their respective positions are $(4, 3, 2, 1, 5)^T$ or $(4, 2, 3, 1, 5)^T$, so that the ranks are given by $(4, 2.5, 2.5, 1, 5)^T$. Spearman's correlation distance between \mathbf{x} and \mathbf{y} is thus given by

$$d_{SC}(\mathbf{x}, \mathbf{y}) = d_{PC}(\mathbf{x}_r, \mathbf{y}_r),$$

which ranges between 0 and 2 and is semi-metric.

The normalized mutual information (NMI) distance is defined as follows. The P elements of \mathbf{x} and \mathbf{y} are considered to be observations of the random variables x and y , respectively. Let $p_x(\cdot)$ denote the PMF of \mathbf{x} found by considering the histogram of the elements of \mathbf{x} with M bins. That is, $p_x(i)$ gives the proportion of the elements $\{x_p\}_{p=1}^P$ in the i^{th} bin. We follow Priness *et al.* (2007) and use the integer value of \sqrt{P} as M . Estimate the entropy of \mathbf{x} by

$$E(x) = - \sum_{m=1}^M p_x(m) \log(p_x(m)),$$

and similarly for \mathbf{y} with PMF $p_y(\cdot)$. Estimate the joint entropy of \mathbf{x} and \mathbf{y} by considering

the joint PMF, denoted $\{p_{xy}(i, j)\}_{i,j=1}^M$, using a two-dimensional histogram of \mathbf{x} and \mathbf{y} , as

$$E(x, y) = - \sum_{m=1}^M \sum_{n=1}^M p_{xy}(m, n) \log(p_{xy}(m, n)).$$

The NMI distance measure is then given by

$$d_{NMI}(\mathbf{x}, \mathbf{y}) = 1 - \frac{E(x) + E(y) - E(x, y)}{\max\{E(x), E(y)\}},$$

where the fraction is the NMI (Michaels *et al.*, 1998). This distance is bounded by 0 and 1 and is semi-metric.

D Simulation results

D.1 Comparison between GRV and Mantel

Since the GRV and Mantel statistics can both be expressed as correlation coefficients, the difference in performance is due to the difference in standardized distance elements used as inputs in each case; a classical standardization is applied to the upper-triangular distance elements by Mantel, and a normalized double-centering is applied to the complete distance matrix by GRV. Figure 2 demonstrates that, for the data simulated as described in Section 3 of the article with $N = 50$ and with the IBS and Mahalanobis distances, the standardization used by GRV yields a stronger correlation than the standardization used by Mantel. This is an example of how the standardization used by GRV better detects the association inherent in the simulated data.

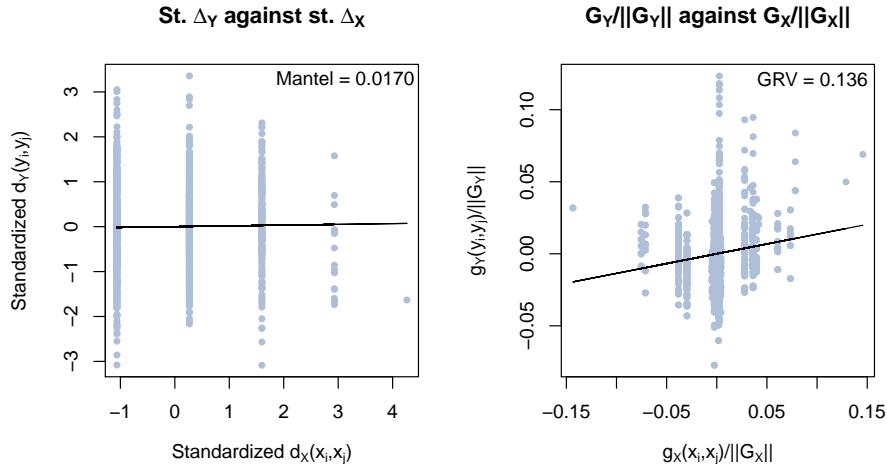


Figure 2: Standardized and normalized double-centered distances (gray points) used by the Mantel and GRV statistics, for the IBS and Mahalanobis distances for $N = 50$. In each case the gradient of the black line equals the correlation and hence the statistic value. The standardization used by GRV yields a higher correlation than achieved by the standardization used by Mantel.

Table 4: Mean (and standard deviation) of the proportion of null p-values less than or equal to the given significance levels for $N = \{30, 50, 70\}$.

N	Significance level (%)		
	1	5	10
30	0.010 (0.007)	0.051 (0.014)	0.100 (0.020)
50	0.010 (0.006)	0.050 (0.016)	0.101 (0.020)
70	0.011 (0.006)	0.051 (0.017)	0.102 (0.021)

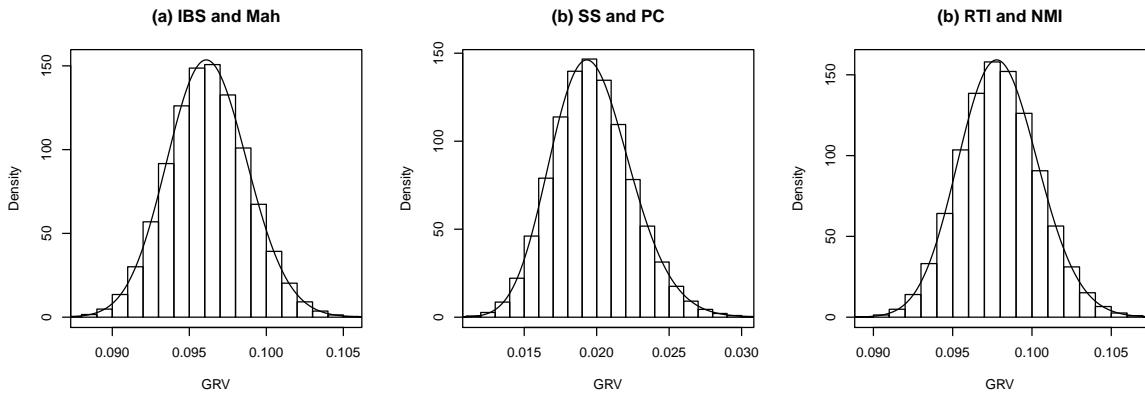


Figure 3: Sampling distributions of the GRV statistic obtained using 10^6 Monte Carlo permutations and the proposed approximate PDF applied to a pathway of the real data including a SNP-set comprised of 2092 SNPs and a transcript-set comprised of 51 transcripts. (a) The IBS distance is applied to the SNP data and the Mahalanobis (Mah) distance is applied to the gene expression data. (b) The SS distance is applied to the SNP data and the Pearson’s correlation distance (PC) distance is applied to the gene expression data. (b) The RTI distance is applied to the SNP data and the normalized mutual information (NMI) distance is applied to the gene expression data.

D.2 Attainment of nominal significance level of GRV

We demonstrate that the GRV test attains the nominal significance level of a given test for a range of significance levels by using a Monte Carlo procedure. We generate paired data inspired by the eQTL application setting used in the power study.

For 100 Monte Carlo runs, 200 paired datasets are generated as described in Section 3 of the article for $P = 2$, $Q = 10$, and $N = \{30, 50, 70\}$, except with $\mathbf{y}_i = \mathbf{e}_i$ so that \mathbf{y}_i is not dependent on \mathbf{x}_i , and the IBS and Mahalanobis distance measures applied. Over all runs the mean and standard deviation of the proportion of p-values less than or equal to significance levels of 1%, 5%, 10% are monitored. These are presented in Table 4, where it is clear that the nominal significance levels are attained.

E Illustration of the null distribution of GRV for real data

In order to illustrate how the null distribution compares with the permutation distribution when applied to real data, we consider a subset of the cancer data. For a SNP-set comprised

of 2092 SNPs and a transcript-set comprised of 51 transcripts, each of the following combinations of SNP and gene expression distances were applied: (a) IBS and Mahalanobis, (b) SS and Pearson's correlation, and (c) RTI and normalized mutual information. For each we obtained the approximate null distribution of the GRV statistic and the permutation distribution using 10^6 Monte Carlo permutations. These are given in Figure 3, where we see that the approximate distribution appears to provide a good fit to the permutation distributions.

References

Agarwal, R. and Kaye, S. B. (2003). Ovarian cancer: strategies for overcoming resistance to chemotherapy. *Nature Reviews Cancer*, **3**, 502–516.

Beckmann, L. *et al.* (2005). Haplotype Sharing Analysis Using Mantel Statistics. *Human Heredity*, **59**(2), 67–78.

Beroukhim, R. *et al.* (2007). Assessing the significance of chromosomal aberrations in cancer: Methodology and application to glioma. *PNAS*, **104**(50), 20007–20012.

Berry, K. J. and Mielke, P. W. (1983). Moment approximations as an alternative to the F test in analysis of variance. *British Journal of Mathematical and Statistical Psychology*, **36**(2), 202–206.

Borg, I. and Groenen, P. J. F. (2005). *Modern multidimensional scaling: theory and applications*. Springer Science+Business Media, Inc., second edition.

Carvalho, B. *et al.* (2007). Exploration, normalization, and genotype calls of high-density oligonucleotide SNP array data. *Biostatistics*, **8**(2), 485–499.

Christopherson, R. I. *et al.* (2002). Inhibitors of de novo nucleotide biosynthesis as drugs. *Accounts of Chemical Research*, **35**(11), 961–971.

Cookson, W. *et al.* (2009). Mapping complex disease traits with global gene expression. *Nature Reviews Genetics*, **10**(3), 184–194.

Dai, W. *et al.* (2011). Systematic CpG Islands Methylation Profiling of Genes in the Wnt Pathway in Epithelial Ovarian Cancer Identifies Biomarkers of Progression-Free Survival. *Clinical Cancer Research*, **17**, 4052.

Dervieux, T. *et al.* (2004). Polyglutamation of methotrexate with common polymorphisms in reduced folate carrier, aminoimidazole carboxamide ribonucleotide transformylase, and thymidylate synthase are associated with methotrexate effects in rheumatoid arthritis. *Arthritis & Rheumatism*, **50**(9), 2766–2774.

Escoufier, Y. (1973). Le traitement des variables vectorielles. *Biometrics*, pages 751–760.

Firestone, S. M. and Davisson, V. J. (1993). A tight binding inhibitor of 5-Aminoimidazole Ribonucleotide Carboxylase. *Journal of Medical Chemistry*, **36**, 3484–3486.

Fong, P. C. *et al.* (2010). Poly (ADP)-ribose polymerase inhibition: frequent durable responses in BRCA carrier ovarian cancer correlating with platinum-free interval. *Journal of Clinical Oncology*, **28**(15), 2512–2519.

Gandy, A. and Rubin-Delanchy, P. (2013). An algorithm to compute the power of monte carlo tests with guaranteed precision. *The Annals of Applied Statistics*, **41**(1), 125–142.

Goh, L. *et al.* (2011). Assessing matched normal and tumor pairs in next-generation sequencing studies. *PloS one*, **6**(3), e17810.

Gower, J. C. (1966). Some distance properties of latent root and vector methods used in multivariate analysis. *Biometrika*, **53**(3-4), 325–338.

Hickey, K. P. o. (1999). Molecular detection of tumour DNA in serum and peritoneal fluid from ovarian cancer patients. *British Journal of Cancer*, **80**(11), 1803–1808.

Hollich, V. *et al.* (2005). Assessment of protein distance measures and tree-building methods for phylogenetic tree reconstruction. *Molecular Biology and Evolution*, **22**(11), 2257–2264.

Irizarry, R. *et al.* (2003). Exploration, normalization, and summaries of high density oligonucleotide array probe level data. *Biostatistics*, **4**(2), 249–264.

Josse, J. *et al.* (2008). Testing the significance of the RV coefficient. *Computational Statistics and Data Analysis*, **53**(1), 82–91.

Jurman, G. *et al.* (2008). Algebraic stability indicators for ranked lists in molecular profiling. *Bioinformatics*, **24**(2), 258–264.

Kamburov, A. *et al.* (2011). ConsensusPathDB: toward a more complete picture of cell biology. *Nucleic Acids Research*, **39**(S1), D712–D717.

Kazi-Aoual, F. *et al.* (1995). Refined approximations to permutation tests for multivariate inference. *Computational statistics and data analysis*, **20**(6), 643–656.

Knijnenburg, T. A. *et al.* (2009). Fewer permutations, more accurate P-values. *Bioinformatics*, **25**(12), i161–i168.

Krzanowski, W. J. (2000). *Principles of multivariate analysis: a user's perspective*. Oxford University Press.

Lasserre, J. (1995). A trace inequality for matrix product. *IEEE Transactions on Automatic Control*, **40**(8), 1500–1501.

Legendre, P. and Fortin, M. J. (2010). Comparison of the Mantel Test and Alternative Approaches for Detecting Complex Multivariate Relationships in the Spatial Analysis of Genetic Data. *Molecular Ecology Resources*, **10**(5), 831–844.

Louhimo, R. and Hautaniemi, S. (2011). CNAMet: an R package for integrating copy number, methylation and expression data. *Bioinformatics*, **27**(6), 887–888.

Mantel, N. (1967). The Detection of Disease Clustering and a Generalized Regression Approach. *Cancer Research*, **27**(2), 209–220.

Mardia, K. *et al.* (1979). *Multivariate analysis*. Academic Press.

Michaels, G. *et al.* (1998). Cluster analysis and data visualization of large-scale gene expression data. In *Pacific Symposium on Biocomputing*, volume 3, pages 42–53.

Mielke, P. W. and Berry, K. J. (2007). *Permutation methods: a distance function approach*. Springer Science+Business Media, LLC.

Minas, C. *et al.* (2011). Distance-based differential analysis of gene curves. *Bioinformatics*, **27**(22), 3135–3141.

Pekalska, E. and Duin, R. P. W. (2005). *The dissimilarity representation for pattern recognition: foundations and applications*. World Scientific Co. Pte. Ltd.

Peres-Neto, P. R. and Jackson, D. A. (2001). How well do multivariate data sets match? The advantages of a Procrustean superimposition approach over the Mantel test. *Oecologia*, **129**(2), 169–178.

Phipson, B. and Smyth, G. K. (2010). Permutation p-values should never be zero: calculating exact p-values when permutations are randomly drawn. *Statistical Applications in Genetics and Molecular Biology*, **9**, 39.

Priness, I. *et al.* (2007). Evaluation of gene-expression clustering via mutual information distance measure. *BMC Bioinformatics*, **8**(1), 111.

Sawyers, C. L. (2008). The cancer biomarker problem. *Nature*, **452**, 548–552.

Selinski, S. and Ickstadt, K. (2005). Similarity Measures for Clustering SNP Data. *Technical Report / Universitt Dortmund, SFB 475 Komplexittsreduktion in Multivariaten Datenstrukturen*.

Shannon, W. D. *et al.* (2002). Mantel statistics to correlate gene expression levels from microarrays with clinical covariates. *Genetic Epidemiology*, **23**(1), 87–96.

Song, L. *et al.* (2012). Comparison of co-expression measures: mutual information, correlation, and model based indices. *BMC Bioinformatics*, **13**(1), 328.

Stronach, E. A. *et al.* (2011). HDAC4-Regulated STAT1 Activation Mediates Platinum Resistance in Ovarian Cancer. *Cancer Research*, **71**, 4412.

Sun, Y. V. *et al.* (2011). Identification of genes associated with complex traits by testing the genetic dissimilarity between individuals. In *BMC Proceedings*, volume 5, page S120. Sun *et al*; licensee BioMed Central Ltd.

Torgerson, W. S. (1952). Multidimensional Scaling: I. Theory and Method. *Psychometrika*, **17**(4), 401–419.

van't Veer, L. J. and Bernards, R. (2008). Enabling personalized cancer medicine through analysis of gene-expression patterns. *Nature*, **452**, 564–570.

Vaughan, S. *et al.* (2011). Rethinking ovarian cancer: recommendations for improving outcomes. *Nature Reviews Cancer*, **11**, 719–725.

Wu, T. J. *et al.* (1997). A measure of DNA sequence dissimilarity based on Mahalanobis distance between frequencies of words. *Biometrics*, **53**, 1431–1439.